SIAM J. SCI. COMPUT.
Vol. 37, No. 5, pp. S581–S601

PROBABILISTIC BOUNDS FOR THE MATRIX CONDITION NUMBER WITH EXTENDED LANCZOS BIDIAGONALIZATION*

SARAH W. GAAF† AND MICHIEL E. HOCHSTENBACH†

Abstract. Reliable estimates for the condition number of a large, sparse, real matrix A are important in many applications. To get an approximation for the condition number $\kappa(A)$, an approximation for the smallest singular value is needed. Standard Krylov subspaces are usually unsuitable for finding a good approximation to the smallest singular value. Therefore, we study extended Krylov subspaces which turn out to be ideal for the simultaneous approximation of both the smallest and largest singular value of a matrix. First, we develop a new extended Lanczos bidiagonalization method. With this method we obtain a lower bound for the condition number. Moreover, the method also yields probabilistic upper bounds for $\kappa(A)$. The user can select the probability with which the upper bound holds, as well as the ratio of the probabilistic upper bound and the lower bound.

Key words. extended Lanczos bidiagonalization, extended Krylov method, matrix condition number, lower bound, probabilistic upper bound

AMS subject classifications. 65F15, 65F35, 65F50, 65F10

DOI. 10.1137/140975218

1. Introduction.

Let $A \in \mathbb{R}^{n\times n}$ be a large, nonsingular matrix. Let $A = X\Sigma Y^T$ be the singular value decomposition of A, where X and Y are $n \times n$ matrices with orthonormal columns containing the left and right singular vectors of A, respectively. Furthermore, Σ is an $n \times n$ diagonal matrix with positive real entries containing the singular values of A that are numbered in decreasing order: $\sigma_1 \geq \cdots \geq \sigma_n > 0$.

We are interested in the important problem of approximating the condition number of A,

$$\kappa(A) = \|A\|\,\|A^{-1}\| = \frac{\sigma_1}{\sigma_n},$$

where $\|\cdot\|$ stands for the 2-norm. The (Golub–Kahan–)Lanczos bidiagonalization method [5] provides an approximation, a lower bound, for the maximum singular value σ_1 of A. In addition, an upper bound for the minimum singular value is obtained, but this is usually a rather poor bound. To approximate the condition number, good approximations to σ_n are needed.

This paper has three contributions. First, we develop a new extended Lanczos bidiagonalization method. The method generates a basis for the extended Krylov subspace:

$$\mathcal{K}^{k+1,k+1}(A^TA,\mathbf{v}) = \mathrm{span}\{(A^TA)^{-k}\mathbf{v},\ldots,(A^TA)^{-1}\mathbf{v},\mathbf{v},A^TA\mathbf{v},\ldots,(A^TA)^k\mathbf{v}\}.$$

Extended Krylov subspace methods have been studied in the last 15 years by various authors [3, 13, 14, 16, 20]. The second contribution of this paper is that we obtain simultaneously a lower bound for σ_1 and an upper bound for σ_n, which leads to a lower bound of good quality for $\kappa(A)$. Third, we obtain a probabilistic upper bound

*Received by the editors June 30, 2014; accepted for publication (in revised form) July 14, 2015; published electronically October 29, 2015.
http://www.siam.org/journals/sisc/37-5/97521.html

†Department of Mathematics and Computer Science, TU Eindhoven, 5600 MB, The Netherlands (s.w.gaaf@tue.nl, m.e.hochstenbach@tue.nl). This work was supported by a Vidi research grant from the Netherlands Organisation for Scientific Research (NWO).

for the condition number. Probabilistic techniques have become increasingly popular; see, for instance, [2, 17, 21, 7, 11]. Whereas in [2, 17, 7] the power method is used, this paper is based on Krylov methods as are the techniques in [17, 21, 11]. An important feature of the Lanczos bidiagonalization procedure is that the starting vector can be (and often is) chosen randomly. Therefore, the probability that this vector has a small component in the direction of the desired singular vector (relative to $1/\sqrt{n}$) is small. Another characteristic of the procedure is that during the bidiagonalization process polynomials implicitly arise. These two properties are exploited in [11] to obtain probabilistic upper bounds for σ_1.

In this paper, we will expand the techniques from [11] to obtain both probabilistic lower bounds for σ_n and probabilistic upper bounds for σ_1, leading to probabilistic upper bounds for $\kappa(A)$. These upper bounds hold with user-chosen probability: the user can select an $\varepsilon > 0$ such that the bounds hold with probability $1 - 2\varepsilon$, as well as a $\zeta > 1$ such that the ratio of the probabilistic upper bound and the lower bound is less than ζ. The method will adaptively perform a number of steps k to accomplish this. Probabilistic condition estimators in [2] or [17] provide a ratio between the probabilistic upper bound and the lower bound, given a fixed k and ε. The method of this paper does not come with an analogous relation; however, the method we propose generally gives sharper bounds as is shown in section 7.

We stress the fact that the method of the present paper requires an (exact) LU decomposition. If this is unaffordable, there are alternative methods available that need only a preconditioner such as an inexact LU decomposition. The Jacobi–Davidson type SVD method [9, 10] is one of these methods. However, because of the current state of both numerical methods and hardware, LU decompositions have increasingly become an option, sometimes also for rather large matrices.

The theory discussed in this paper considers only real matrices. For general complex matrices the theory from this paper to obtain probabilistic bounds needs to be adapted in a nontrivial way, and will be subject to future study.

The rest of this paper is organized as follows. In section 2 we introduce the extended Lanczos bidiagonalization method, and the special structure of the matrices obtained by this method are examined in section 3. Section 4 focuses on the Laurent polynomials arising in the procedure. In section 5 we elaborate on the computation of a probabilistic bound for the condition number. Section 6 discusses some comparisons with several other (probabilistic) condition number estimators. We end with some numerical experiments and conclusions in sections 7 and 8.

2. Extended Lanczos bidiagonalization.

The method we will develop starts with a random vector $\mathbf{v}_0$ with unit norm. We express $\mathbf{v}_0$ as a linear combination of the right singular vectors $\mathbf{y}_i$ of A,

$$\mathbf{v}_0 = \sum_{i=1}^{n} \gamma_i \, \mathbf{y}_i. \tag{2.1}$$

Notice that both the $\mathbf{y}_i$ and γ_i are unknown. The extended Lanczos bidiagonalization method repeatedly applies the matrices A, A^T, A^{-T}, and A^{-1}. In every step a generated vector is orthogonalized with respect to the previously constructed vectors, and subsequently normalized. This procedure can be visualized as a string of operations working on vectors:

$$\mathbf{v}_0 \xrightarrow{A} \mathbf{u}_0 \xrightarrow{A^T} \mathbf{v}_1 \xrightarrow{A^{-T}} \mathbf{u}_{-1} \xrightarrow{A^{-1}} \mathbf{v}_{-1} \xrightarrow{A} \mathbf{u}_1 \xrightarrow{A^T} \dots .$$

Note that in this visualization the orthonormalization of the vectors is not shown. In this scheme, applying the operation A^{-T} after A^T (and A after A^{-1}) may seem contradictory, but since the vectors are orthogonalized in between this truly yields new vectors. Another way to represent this procedure is the table below:

Step	Action	Generated	Action	Generated	Action	Generated	Action	Generated
0	$A\mathbf{v}_0$	$\mathbf{u}_0$	$A^T\mathbf{u}_0$	$\mathbf{v}_1$	$A^{-T}\mathbf{v}_1$	$\mathbf{u}_{-1}$	$A^{-1}\mathbf{u}_{-1}$	$\mathbf{v}_{-1}$
1	$A\mathbf{v}_{-1}$	$\mathbf{u}_1$	$A^T\mathbf{u}_1$	$\mathbf{v}_2$	$A^{-T}\mathbf{v}_2$	$\mathbf{u}_{-2}$	$A^{-1}\mathbf{u}_{-2}$	$\mathbf{v}_{-2}$
$\vdots$								
$k-1$	$A\mathbf{v}_{-k+1}$	$\mathbf{u}_{k-1}$	$A^T\mathbf{u}_{k-1}$	$\mathbf{v}_k$	$A^{-T}\mathbf{v}_k$	$\mathbf{u}_{-k}$	$A^{-1}\mathbf{u}_{-k}$	$\mathbf{v}_{-k}$

During the procedure, the generated vectors $\mathbf{v}_j$ are normalized after being orthogonalized with respect to all previously generated $\mathbf{v}_i$, i.e., for $k \geq 1$

$$\mathbf{v}_k \perp \{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_{k-1}, \mathbf{v}_{-k+1}\}, \qquad \mathbf{v}_{-k} \perp \{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_{-k+1}, \mathbf{v}_k\}.$$

Similarly, all generated vectors $\mathbf{u}_j$ have unit norm and

$$\mathbf{u}_{k-1} \perp \{\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{k-2}, \mathbf{u}_{-k+1}\}, \qquad \mathbf{u}_{-k} \perp \{\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{-k+1}, \mathbf{u}_{k-1}\}.$$

Define the matrices $V_1 = [\mathbf{v}_0]$ and $U_1 = [\mathbf{u}_0]$, and for $k \geq 1$

$$\begin{aligned} V_{2k} &= [V_{2k-1}, \mathbf{v}_k], & U_{2k} &= [U_{2k-1}, \mathbf{u}_{-k}], \\ V_{2k+1} &= [V_{2k}, \mathbf{v}_{-k}], & U_{2k+1} &= [U_{2k}, \mathbf{u}_k]. \end{aligned}$$

The columns of these matrices are orthonormal and span the corresponding subspaces $\mathcal{V}_{2k}$, $\mathcal{V}_{2k+1}$, $\mathcal{U}_{2k}$, and $\mathcal{U}_{2k+1}$, respectively. We assume for the moment that no breakdowns occur, so all spaces are of full dimension; how to handle a breakdown is discussed in section 7. After $k \geq 1$ steps the algorithm gives rise to the following matrix equations:

$$\begin{aligned} AV_{2k-1} &= U_{2k-1}H_{2k-1}, \\ A^TU_{2k-1} &= V_{2k}(H_{2k-1,2k})^T &&= V_{2k-1}(H_{2k-1})^T + \beta_{k-1}\,\mathbf{v}_k\,\mathbf{e}_{2k-1}^T, \\ A^{-T}V_{2k} &= U_{2k}(K_{2k})^T, \\ A^{-1}U_{2k} &= V_{2k+1}K_{2k+1,2k} &&= V_{2k}K_{2k} + \delta_k\,\mathbf{v}_{-k}\,\mathbf{e}_{2k}^T. \end{aligned} \tag{2.2}$$

$$\begin{aligned} AV_{2k} &= U_{2k+1}H_{2k+1,2k} &&= U_{2k}H_{2k} + \beta_{-k}\,\mathbf{u}_k\,\mathbf{e}_{2k}^T, \\ A^TU_{2k} &= V_{2k}(H_{2k})^T, \\ A^{-T}V_{2k+1} &= U_{2k+2}(K_{2k+1,2k+2})^T &&= U_{2k+1}(K_{2k+1})^T + \delta_{-k}\,\mathbf{u}_{-k-1}\,\mathbf{e}_{2k+1}^T, \\ A^{-1}U_{2k+1} &= V_{2k+1}K_{2k+1}. \end{aligned} \tag{2.3}$$

Here, and throughout this paper, $H_{m,p}$ is an $m \times p$ matrix. We will use only one subscript if the matrix is square, i.e., H_m is an $m \times m$ matrix, and we will refer to the matrices $H_{m,p}$ and $K_{m,p}$ as H and K if the size is not of interest. Furthermore, $\mathbf{e}_i$ is the ith unit vector and the coefficients β_j and δ_j are entries of the matrices H and K, which will be specified in section 3. More details on the recurrence relation between the vectors $\mathbf{u}$ and $\mathbf{v}$ will be given in (3.3) where we show that orthogonalization can be done using 3-term recurrences. In particular, the pseudocode for the algorithm that will be introduced in section 7 shows that only three vectors of storage are needed.

Let $\theta_1^{(2k-1)} \geq \cdots \geq \theta_{2k-1}^{(2k-1)}$ be the singular values of H_{2k-1}, and let $\theta_1^{(2k)} \geq \cdots \geq \theta_{2k}^{(2k)}$ be the singular values of H_{2k}. Similarly, let $\xi_1^{(2k-1)} \geq \cdots \geq \xi_{2k-1}^{(2k-1)}$ be the singular values of K_{2k-1}, and let $\xi_1^{(2k)} \geq \cdots \geq \xi_{2k}^{(2k)}$ be the singular values of K_{2k}. These values are approximations of the singular values of A and A^{-1}, respectively. We will avoid the use of superscripts if this is clear from the context. Further, let $\mathbf{c}_j$ and $\mathbf{d}_j$ indicate the corresponding right singular vectors of H and K, respectively. We will now study the behavior of these values θ_j and ξ_j to obtain bounds for the extreme singular values of A.

PROPOSITION 2.1.

(a) *For $1 \leq j \leq 2k-1$ the singular values of H converge monotonically to the largest singular values of A: $\theta_j^{(2k-1)} \leq \theta_j^{(2k)} \leq \sigma_j(A)$.*

(b) *For $1 \leq j \leq 2k-1$ the inverse singular values of K converge monotonically to the smallest singular values of A:*

$$\sigma_{n-j+1}(A) = \left(\sigma_j(A^{-1})\right)^{-1} \leq \left(\xi_j^{(2k)}\right)^{-1} \leq \left(\xi_j^{(2k-1)}\right)^{-1}.$$

Proof. The matrix H_{2k-1} can be seen as the matrix H_{2k} from which the $2k$th row and column have been deleted. The same holds for the matrices K_{2k-1} and K_{2k}. Now we apply [12, Cor. 3.1.3] and obtain the first inequalities of both (a) and (b). The second inequalities hold because of [12, Lem. 3.3.1]. ∎

In the next section we will see that $H^{-1} = K$, which means that $\{\theta_1^{-1}, \ldots, \theta_{2k}^{-1}\} = \{\xi_1, \ldots, \xi_{2k}\}$. Proposition 2.1 shows in particular that the largest singular value of the matrices H converges monotonically to σ_1, and the inverse of the largest singular value of the matrices K converges monotonically to σ_n. After the kth step of the procedure, we obtain the value $\theta_1^{(2k)}$, a lower bound for σ_1, and the value $(\xi_1^{(2k)})^{-1}$, an upper bound for σ_n.

COROLLARY 2.2. *After the kth step of extended Lanczos bidiagonalization we obtain a lower bound for the condition number of A:*

$$\kappa_{\text{low}}(A) = \frac{\theta_1}{\xi_1^{-1}} \leq \frac{\sigma_1}{\sigma_n} = \kappa(A). \tag{2.4}$$

The experiments in section 7 show for different matrices that the lower bound achieved by extended Lanczos bidiagonalization may often be very good.

We can reformulate the expressions in (2.2) and (2.3) to see the similarities with the extended Lanczos method (see, e.g., [13]) with starting vector $\mathbf{v}_0$ and matrix A^TA, so that for $k \geq 1$:

$$\begin{aligned}
A^TAV_{2k-1} &= A^TU_{2k-1}H_{2k-1} \\
&= V_{2k-1}(H_{2k-1})^TH_{2k-1} + \beta_{k-1}\,\mathbf{v}_k\,\mathbf{e}_{2k-1}^T\,H_{2k-1}, \\
(A^TA)^{-1}V_{2k} &= A^{-1}U_{2k}(K_{2k})^T \\
&= V_{2k}K_{2k}(K_{2k})^T + \alpha_k^{-1}\,\delta_k\,\mathbf{v}_{-k}\,\mathbf{e}_{2k}^T, \\
AA^TU_{2k} &= AV_{2k}(H_{2k})^T \\
&= U_{2k}H_{2k}(H_{2k})^T + \alpha_k\,\beta_{-k}\,\mathbf{u}_k\,\mathbf{e}_{2k}^T, \\
(AA^T)^{-1}U_{2k-1} &= A^{-T}V_{2k-1}K_{2k-1} \\
&= U_{2k-1}(K_{2k-1})^TK_{2k-1} + \delta_{-k+1}\,\mathbf{u}_{-k+1}\,\mathbf{e}_{2k-1}^TK_{2k-1}.
\end{aligned} \tag{2.5}$$

This way of representing the procedure will be convenient in the next sections where we will investigate the structure of the generated matrices and introduce Laurent polynomials.

3. The special structure of the generated H and K matrices.

In the previous section we introduced the extended Lanczos bidiagonalization method. The four leading submatrices arising in (2.2) and (2.3) are given by

$$\begin{aligned} H_{2k-1} &= U_{2k-1}^T A V_{2k-1}, & K_{2k-1} &= V_{2k-1}^T A^{-1} U_{2k-1}, \\ H_{2k} &= U_{2k}^T A V_{2k}, & K_{2k} &= V_{2k}^T A^{-1} U_{2k}. \end{aligned}$$

These matrices H and K turn out to be tridiagonal matrices with a special structure as we will show in the next proposition. Note that we assume for all $j \in \{-k, \dots, k\}$ that the entries α_j, β_j, and δ_j are nonzero.

PROPOSITION 3.1.

(a) *The matrix H is tridiagonal and of the form*

$$\begin{bmatrix} \alpha_0 & \beta_0 & & & & & \\ & \alpha_1 & & & & & \\ & \beta_{-1} & \alpha_{-1} & \beta_1 & & & \\ & & & \alpha_2 & & & \\ & & & \beta_{-2} & \alpha_{-2} & \beta_2 & \\ & & & & & \alpha_3 & \\ & & & & & & \ddots \end{bmatrix}, \tag{3.1}$$

where its entries satisfy

$$\begin{aligned} h_{2j,2j} &= \alpha_j &&= \|A^{-T}\mathbf{v}_j\|^{-1} = \|A^T\mathbf{u}_{-j}\|, \\ h_{2j+1,2j} &= \beta_{-j} &&= \mathbf{u}_j^T A \mathbf{v}_j, \\ h_{2j+1,2j+1} &= \alpha_{-j} &&= \mathbf{u}_j^T A \mathbf{v}_{-j}, \\ h_{2j+1,2j+2} &= \beta_j &&= \|A^T\mathbf{u}_j - (\mathbf{u}_j^T A\mathbf{v}_j)\mathbf{v}_j - (\mathbf{u}_j^T A \mathbf{v}_{-j})\mathbf{v}_{-j}\|. \end{aligned}$$

(b) *The matrix K is tridiagonal and of the form*

$$\begin{bmatrix} \alpha_0^{-1} & \delta_0 & & & & & \\ & \alpha_1^{-1} & & & & & \\ & \delta_1 & \alpha_{-1}^{-1} & \delta_{-1} & & & \\ & & & \alpha_2^{-1} & & & \\ & & & \delta_2 & \alpha_{-2}^{-1} & \delta_{-2} & \\ & & & & & \alpha_3^{-1} & \\ & & & & & & \ddots \end{bmatrix}, \tag{3.2}$$

where its diagonal entries are defined in (a) *and its off-diagonal entries satisfy*

$$\begin{aligned} k_{2j+1,2j} &= \delta_j &&= \|A^{-1}\mathbf{u}_{-j} - (\mathbf{v}_{-j+1}^T A^{-1}\mathbf{u}_{-j})\mathbf{v}_{-j+1} - (\mathbf{v}_j^T A^{-1}\mathbf{u}_{-j})\mathbf{v}_j\|, \\ k_{2j+1,2j+2} &= \delta_{-j} &&= \mathbf{v}_{-j}^T A^{-1}\mathbf{u}_{-(j+1)}. \end{aligned}$$

Proof. We will focus first on the transposed matrix H_{2k}^T. Note that $A^T U_{2k} = V_{2k} H_{2k}^T$. The $(2j+1)$st column of $A^T U_{2k}$ is $A^T \mathbf{u}_j$, and thus the $(2j+1)$st column of H_{2k}^T, can be described using the step of the algorithm in which the vector $\mathbf{v}_{j+1}$ is constructed:

$$\beta_j \mathbf{v}_{j+1} = A^T \mathbf{u}_j - \sum_{i=-j}^{j} \gamma_i \, \mathbf{v}_i,$$

where $\gamma_i = \mathbf{v}_i^T A^T \mathbf{u}_j = \mathbf{u}_j^T A \mathbf{v}_i$ and β_j is a factor such that $\mathbf{v}_{j+1}$ has unit norm. For all $i \in \{-j+1, \ldots, j-1\}$ we have

$$\begin{aligned} A\mathbf{v}_i &\in \operatorname{span}\{(AA^T)^{-j+1} A\mathbf{v}_0, \, \ldots, \, (AA^T)^{j-1} A\mathbf{v}_0\} \\ &= \operatorname{span}\{\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{-j+1}, \mathbf{u}_{j-1}\}, \end{aligned}$$

and therefore $\gamma_i = 0$ for all $i \in \{-j+1, \ldots, j-1\}$. We obtain the three-term recurrence relation

$$A^T \mathbf{u}_j = (\mathbf{u}_j^T A \mathbf{v}_j) \, \mathbf{v}_j + (\mathbf{u}_j^T A \mathbf{v}_{-j}) \, \mathbf{v}_{-j} + \beta_j \mathbf{v}_{j+1},$$

which implies that the $(2j+1)$st column of H_{2k}^T has only three nonzero entries. This gives us the three nonzero entries of the odd rows of H: $h_{2j+1,2j}$, $h_{2j+1,2j+1}$, and $h_{2j+1,2j+2}$.

For the description of the $(2j)$th column of H_{2k}^T, another step of the algorithm is used, namely

$$\alpha_j^{-1} \mathbf{u}_{-j} = A^{-T} \mathbf{v}_j - \sum_{i=-j+1}^{j-1} \gamma_i \mathbf{u}_i,$$

where $\gamma_i = \mathbf{u}_i^T A^{-T} \mathbf{v}_j = \mathbf{v}_j^T A^{-1} \mathbf{u}_i$ and α_j^{-1} is a factor such that $\mathbf{u}_{-j}$ has unit norm. For all $i \in \{-j+1, \ldots, j-1\}$ we have

$$\begin{aligned} A^{-1}\mathbf{u}_i &\in \operatorname{span}\{(A^T A)^{-j+1} \mathbf{v}_0, \, \ldots, \, (A^T A)^{j-1} \mathbf{v}_0\} \\ &= \operatorname{span}\{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_{j-1}, \mathbf{v}_{-j+1}\}, \end{aligned}$$

and therefore $\gamma_i = 0$ for all $i \in \{-j+1, \ldots, j-1\}$. We obtain the recurrence relation

$$A^{-T} \mathbf{v}_j = \alpha_j^{-1} \mathbf{u}_{-j}, \qquad \text{and therefore} \qquad A^T \mathbf{u}_{-j} = \alpha_j \mathbf{v}_j,$$

implying that the $(2j)$th column of H_{2k}^T has only one nonzero entry. The entries of the matrix K can be obtained by a similar reasoning. $\square$

This description of the matrices H and K leads to the following recurrence relations:

$$\begin{array}{llll} A\mathbf{v}_{-k} & = \alpha_{-k} \mathbf{u}_k, & k \geq 0, & \star \\ A\mathbf{v}_k & = \beta_{k-1} \mathbf{u}_{k-1} + \alpha_k \mathbf{u}_{-k} + \beta_{-k} \mathbf{u}_k, & k \geq 1, & \\ A^T \mathbf{u}_{-k} & = \alpha_k \mathbf{v}_k, & k \geq 1, & \\ A^T \mathbf{u}_k & = \beta_{-k} \mathbf{v}_k + \alpha_{-k} \mathbf{v}_{-k} + \beta_k \mathbf{v}_{k+1}, & k \geq 1, & \star \\ A^{-T} \mathbf{v}_k & = \alpha_k^{-1} \mathbf{u}_{-k}, & k \geq 1, & \star \\ A^{-T} \mathbf{v}_{-k} & = \delta_k \mathbf{u}_{-k} + \alpha_{-k}^{-1} \mathbf{u}_k + \delta_{-k} \mathbf{u}_{-(k+1)}, & k \geq 1, & \\ A^{-1} \mathbf{u}_k & = \alpha_{-k}^{-1} \mathbf{v}_{-k}, & k \geq 0, & \\ A^{-1} \mathbf{u}_{-k} & = \delta_{-(k-1)} \mathbf{v}_{-(k-1)} + \alpha_k^{-1} \mathbf{v}_k + \delta_k \mathbf{v}_{-k}, & k \geq 1, & \star \end{array} \tag{3.3}$$

and $A^T\mathbf{u}_0 = \alpha_0\mathbf{v}_0 + \beta_0\mathbf{v}_1$, $A^{-T}\mathbf{v}_0 = \alpha_0^{-1}\mathbf{u}_0 + \delta_0\mathbf{u}_{-1}$. The relations indicated by a $\star$ correspond to the matrix vector multiplications that are done explicitly during the procedure, while the other lines are added to give a complete representation of the relations in (2.2) and (2.3). These relations suggest that this method requires at most six vectors of storage, and the algorithm presented in section 7 even shows only three vectors have to be stored. Furthermore, having found this explicit form of the two matrices, it can be seen that the matrices H and K are inverses.

PROPOSITION 3.2. *The leading submatrix of H of order j is the inverse of the leading submatrix of K of the same order, i.e., for $1 \leq j < n$,*

$$H_jK_j = K_jH_j = I_j.$$

Proof. If we would carry out n steps of extended Lanczos bidiagonalization, we would obtain orthogonal matrices V_n and U_n satisfying

$$\begin{aligned} H_nK_n &= U_n^TAV_nV_n^TA^{-1}U_n = I_n, \\ K_nH_n &= V_n^TA^{-1}U_nU_n^TAV_n = I_n. \end{aligned}$$

Due to the special tridiagonal structure, it is easy to see that the statement of the proposition holds. ∎

The previous proposition implies that the singular values of K are the inverses of the singular values of H, and therefore we can adjust Corollary 2.2.

COROLLARY 3.3. *After the kth step of extended Lanczos bidiagonalization we obtain a lower bound for the condition number of A:*

$$\kappa_{\text{low}}(A) = \frac{\theta_1}{\theta_{2k}} \leq \frac{\sigma_1}{\sigma_n} = \kappa(A). \tag{3.4}$$

The matrices in the reformulated expressions (2.5) also have a special structure, just as the matrices formed in the extended Lanczos method in [13]. The four symmetric matrices generated in this extended Lanczos process, for $k \geq 1$, are given by

$$\begin{aligned} R_{2k-1} &= (H_{2k-1})^TH_{2k-1} &&= V_{2k-1}^TA^TAV_{2k-1}, \\ \widetilde{R}_{2k} &= H_{2k}H_{2k}^T &&= U_{2k}^TAA^TU_{2k}, \\ \widetilde{S}_{2k-1} &= (K_{2k-1})^TK_{2k-1} &&= U_{2k-1}^T(AA^T)^{-1}U_{2k-1}, \\ S_{2k} &= K_{2k}K_{2k}^T &&= V_{2k}^T(A^TA)^{-1}V_{2k}. \end{aligned} \tag{3.5}$$

All four are the product of two tridiagonal matrices with a special structure, namely the matrices obtained from extended Lanczos bidiagonalization. The matrices R_{2k-1} and $\widetilde{R}_{2k}$ are pentadiagonal and of the form

$$R_{2k-1} = \begin{bmatrix} \times & \times & & & & & \\ \times & \times & \times & \times & & & \\ & \times & \times & \times & & & \\ & \times & \times & \times & \times & \times & \\ & & & \times & \times & \times & \\ & & & \times & \times & \ddots & \end{bmatrix}, \quad \widetilde{R}_{2k} = \begin{bmatrix} \times & \times & \times & & & & \\ \times & \times & \times & & & & \\ \times & \times & \times & \times & \times & & \\ & & \times & \times & \times & & \\ & & \times & \times & \times & \times & \\ & & & & \times & \ddots & \end{bmatrix}.$$

The matrices S_{2k} and $\widetilde{S}_{2k+1}$ have similar structures. The product of the matrices R and S is a rank-one modification of the identity. Again, if we would carry out n steps of extended Lanczos bidiagonalization, we would obtain orthogonal matrices V_n and U_n with the following identities:

$$\begin{array}{rclclcl} S_nR_n &=& V_n^T(A^TA)^{-1}V_nV_n^T(A^TA)V_n &=& I_n &=& R_nS_n, \\ \widetilde{S}_n\widetilde{R}_n &=& U_n^T(AA^T)^{-1}U_nU_n^T(AA^T)U_n &=& I_n &=& \widetilde{R}_n\widetilde{S}_n. \end{array}$$

Due to the special pentadiagonal structure of the matrices, for $1 \le j < n$ the product is a rank-one modification of the identity, where we have to distinguish between the even and odd cases:

$$\begin{array}{ll} S_{2k}R_{2k} = I_{2k} + \mathring{\mathbf{w}}_{2k}\mathbf{e}_{2k}^T, & S_{2k+1}R_{2k+1} = I_{2k+1} + \mathbf{e}_{2k+1}\mathring{\mathbf{w}}_{2k+1}^T, \\ R_{2k}S_{2k} = I_{2k} + \mathbf{e}_{2k}\check{\mathbf{w}}_{2k}^T, & R_{2k+1}S_{2k+1} = I_{2k+1} + \check{\mathbf{w}}_{2k+1}\mathbf{e}_{2k+1}^T, \\ \widetilde{S}_{2k}\widetilde{R}_{2k} = I_{2k} + \widetilde{\mathbf{w}}_{2k}\mathbf{e}_{2k}^T, & \widetilde{R}_{2k+1}\widetilde{S}_{2k+1} = I_{2k+1} + \mathbf{e}_{2k+1}\widetilde{\mathbf{w}}_{2k+1}^T, \\ \widetilde{R}_{2k}\widetilde{S}_{2k} = I_{2k} + \mathbf{e}_{2k}\widehat{\mathbf{w}}_{2k}^T, & \widetilde{S}_{2k+1}\widetilde{R}_{2k+1} = I_{2k+1} + \widehat{\mathbf{w}}_{2k+1}\mathbf{e}_{2k+1}^T. \end{array}$$

Here, the various vectors $\mathbf{w}_{2k} \in \mathbb{R}^{2k}$ and $\mathbf{w}_{2k+1} \in \mathbb{R}^{2k+1}$ are such that only the last two entries are (possibly) nonvanishing.

The matrices S, $\widetilde{S}$, R, and $\widetilde{R}$ are used in the next section to give an explicit expression for the Laurent polynomials arising in extended Lanczos bidiagonalization.

4. Polynomials arising in extended Lanczos bidiagonalization.

In every step of the extended Lanczos bidiagonalization procedure four different vectors are generated. Since these vectors lie in an extended Krylov subspace, they can be expressed using polynomials:

$$\begin{array}{rclcl} \mathbf{v}_k &=& p_k(A^TA)\mathbf{v}_0 &\in& \mathcal{K}^{k,k+1}(A^TA, \mathbf{v}_0), \\ \mathbf{u}_{-k} &=& q_{-k}(AA^T)A\mathbf{v}_0 &\in& \mathcal{K}^{k+1,k}(AA^T, A\mathbf{v}_0), \\ \mathbf{v}_{-k} &=& p_{-k}(A^TA)\mathbf{v}_0 &\in& \mathcal{K}^{k+1,k+1}(A^TA, \mathbf{v}_0), \\ \mathbf{u}_k &=& q_k(AA^T)A\mathbf{v}_0 &\in& \mathcal{K}^{k+1,k+1}(AA^T, A\mathbf{v}_0). \end{array} \tag{4.1}$$

The polynomials p_k and p_{-k} are Laurent polynomials of the form

$$p_k(t) = \sum_{j=-k+1}^{k} \mu_j^{(k)}\, t^j, \qquad p_{-k}(t) = \sum_{j=-k}^{k} \mu_j^{(-k)}\, t^j. \tag{4.2}$$

Similarly, q_{-k} and q_k are Laurent polynomials and are defined as

$$q_{-k}(t) = \sum_{j=-k}^{k-1} \nu_j^{(-k)}\, t^j, \qquad q_k(t) = \sum_{j=-k}^{k} \nu_j^{(k)}\, t^j. \tag{4.3}$$

The recurrence relations in (3.3) give rise to recurrence relations connecting the polynomials p and q:

$$\begin{array}{rcll} p_{-k}(t) &=& \alpha_{-k}q_k(t), & k \ge 0, \\ p_k(t) &=& \beta_{k-1}q_{k-1}(t) + \alpha_kq_{-k}(t) + \beta_kq_k(t), & k \ge 1, \\ tq_{-k}(t) &=& \alpha_kp_k(t), & k \ge 1, \\ tq_k(t) &=& \beta_{-k}p_k(t) + \alpha_{-k}p_{-k}(t) + \beta_kp_{k+1}(t), & k \ge 1, \\ t^{-1}p_k(t) &=& \alpha_k^{-1}q_{-k}(t), & k \ge 1, \\ t^{-1}p_{-k}(t) &=& \delta_kq_{-k}(t) + \alpha_{-k}^{-1}q_k(t) + \delta_{-k}q_{-(k+1)}(t), & k \ge 1, \\ q_k(t) &=& \alpha_{-k}^{-1}p_{-k}(t), & k \ge 0, \\ q_{-k}(t) &=& \delta_{-(k-1)}p_{-(k-1)}(t) + \alpha_k^{-1}p_k(t) + \delta_kp_{-k}(t), & k \ge 1, \end{array}$$

and $tq_0(t) = \alpha_0 p_0(t) + \beta_0 p_1(t)$, $t^{-1}p_0(t) = \alpha_0^{-1} q_0(t) + \delta_0 q_{-1}(t)$.

Define the following two inner products:

$$\langle f, g\rangle = \mathbf{v}_0^T f(A^TA)\, g(A^TA)\mathbf{v}_0, \tag{4.4}$$

$$[f, g] = \mathbf{v}_0^T f(A^TA)\, A^TA\, g(A^TA)\mathbf{v}_0. \tag{4.5}$$

LEMMA 4.1. *Let $i, j \in \{-k, \dots, k\}$. The polynomials p_i and p_j are orthonormal with respect to the inner product* (4.4)*, whilst the polynomials q_i and q_j are orthonormal with respect to the inner product* (4.5)*.*

Proof. By construction of the $\mathbf{v}_i$'s and $\mathbf{u}_i$'s we have

$$\langle p_i, p_j\rangle = \mathbf{v}_0^T p_i(A^TA)\, p_j(A^TA)\mathbf{v}_0 = \mathbf{v}_i^T\mathbf{v}_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j \end{cases}$$

and

$$[q_i, q_j] = \mathbf{v}_0^T q_i(A^TA)\, A^TA\, q_j(A^TA)\mathbf{v}_0 = \mathbf{u}_i^T\mathbf{u}_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases} \qquad \square$$

Recall that, for $1 < j \le 2k$, θ_j is a singular value of H, ξ_j is a singular value of K, and $\mathbf{c}_j$ and $\mathbf{d}_j$ indicate the corresponding right singular vectors of H and K, respectively.

PROPOSITION 4.2.

(a) *The zeros of the polynomial p_k are exactly $\theta_1^2, \dots, \theta_{2k-1}^2$.*
(b) *The zeros of the polynomial p_{-k} are exactly $\theta_1^2, \dots, \theta_{2k}^2$.*
(c) *The zeros of the polynomial q_{-k} are exactly $\theta_1^2, \dots, \theta_{2k-1}^2$.*
(d) *The zeros of the polynomial q_k are exactly $\theta_1^2, \dots, \theta_{2k}^2$.*

Proof. The proof is similar for all of the polynomials; we will only give details for the first two. Starting with p_k, let $j \in \{1, \dots, 2k-1\}$. Using (2.5) it can be easily seen that the Galerkin condition holds for the pair $(\theta_j^2, V_{2k-1}\mathbf{c}_j)$:

$$A^TAV_{2k-1}\mathbf{c}_j - \theta_j^2\, V_{2k-1}\mathbf{c}_j \perp \mathcal{V}_{2k-1}.$$

Further, since $V_{2k-1}\mathbf{c}_j \in \mathcal{V}_{2k-1}$ it follows that

$$(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j \in \operatorname{span}\{(A^TA)^{-k+1}\mathbf{v}_0, \dots, (A^TA)^k\mathbf{v}_0\}.$$

For each $j = 1, \dots, 2k-1$ we have that $(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j \in \mathcal{V}_{2k}$ but is orthogonal to $\mathcal{V}_{2k-1}$. This means that for all $j = 1, \dots, 2k-1$ the vector $(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j$ is a nonzero multiple of $\mathbf{v}_k = p_k(A^TA)\mathbf{v}_0$. Hence $p_k(t)$ contains all factors $t - \theta_j^2$, i.e., its zeros are exactly $\theta_1^2, \dots, \theta_{2k-1}^2$.

Similarly for the polynomial p_{-k}, let $i \in \{1, \dots, 2k\}$. Again, using (2.5), it can be easily seen that the Galerkin condition holds for the pair $(\xi_i^2, V_{2k}\mathbf{d}_i)$. For each $i = 1, \dots, 2k$ the vector $((A^TA)^{-1} - \xi_i^2 I)V_{2k}\mathbf{d}_i$ is a nonzero multiple of $\mathbf{v}_{-k} = p_{-k}(A^TA)\mathbf{v}_0$, since it is orthogonal to $\mathcal{V}_{2k}$ but an element of $\mathcal{V}_{2k+1}$. Thus p_{-k} contains all factors $(t^{-1} - \xi_j^2)$, and thus all of the factors $(t^{-1} - \theta_j^{-2})$, since $H^{-1} = K$.

Similar proofs can be given for (c) and (d). Note that the proofs in [11, p. 467] and [19, pp. 266–267] follow the same line of reasoning. $\square$

We know from (4.2) that $p_k(t) = \sum_{j=-k+1}^{k} \mu_j^{(k)} t^j$, which, using the result of Proposition 4.2, implies that p_k is of the form

$$p_k(t) = \mu_k^{(k)} \cdot t^{-k+1} \cdot (t - \theta_1^2) \cdots (t - \theta_{2k-1}^2). \tag{4.6}$$

Similarly, p_{-k}, q_k, and q_{-k} are of the form

$$
\begin{aligned}
p_{-k}(t) &= \mu_{-k}^{(-k)} \cdot t^k \cdot (t^{-1} - \theta_1^{-2}) \cdots (t^{-1} - \theta_{2k}^{-2}), \\
q_{-k}(t) &= \nu_{-k}^{(-k)} \cdot t^{k-1} \cdot (t^{-1} - \theta_1^{-2}) \cdots (t^{-1} - \theta_{2k-1}^{-2}), \\
q_k(t) &= \nu_k^{(k)} \cdot t^{-k+1} \cdot (t - \theta_1^2) \cdots (t - \theta_{2k}^2).
\end{aligned}
\tag{4.7}
$$

It turns out that the coefficients $\mu_k^{(k)}$, $\mu_{-k}^{(-k)}$, $\nu_{-k}^{(-k)}$, and $\nu_k^{(k)}$ can be expressed as a product of certain entries of the matrices H and K introduced in (3.1) and (3.2), respectively.

LEMMA 4.3. *The coefficients $\mu_k^{(k)}$, $\mu_{-k}^{(-k)}$, $\nu_{-k}^{(-k)}$, and $\nu_k^{(k)}$ of the polynomials p_k, p_{-k}, q_{-k}, and q_k can be expressed as the product of entries of the matrices H and K defined in* (3.1) *and* (3.2)*, respectively:*

$$
\begin{aligned}
\nu_{-k}^{(-k)} &= (-1)^k \prod_{i=-k+1}^{k} \alpha_i \prod_{i=0}^{k-1} \beta_i^{-1} \prod_{i=1}^{k-1} \delta_i^{-1} \quad \text{and} \quad \mu_{-k}^{(-k)} = \delta_k^{-1} \nu_{-k}^{(-k)}, \\
\nu_k^{(k)} &= (-1)^k \prod_{i=-k}^{k} \alpha_i^{-1} \prod_{i=0}^{k-1} \beta_i^{-1} \prod_{i=1}^{k} \delta_i^{-1} \quad \text{and} \quad \mu_k^{(k)} = \beta_{k-1}^{-1} \nu_{k-1}^{(k-1)}.
\end{aligned}
\tag{4.8}
$$

Proof. From the equations in (2.2), the expressions in (4.1), and from the form of the matrices H and K whose entries are described explicitly in (3.1) and (3.2), respectively, the following recurrence relations for the polynomials can be derived:

$$
\begin{aligned}
q_{-k}(t) &= \alpha_k t^{-1} p_k(t), \\
p_{-k}(t) &= \delta_k^{-1}(q_{-k}(t) - \delta_{-k+1} p_{-k+1}(t) - \alpha_k^{-1} p_k(t)), \\
q_k(t) &= \alpha_{-k}^{-1} p_{-k}(t), \\
p_{k+1}(t) &= \beta_k^{-1}(t q_k(t) - \beta_k p_k(t) - \alpha_{-k} p_{-k}(t)).
\end{aligned}
\tag{4.9}
$$

Manipulating these relations one obtains recurrence relations for the coefficients:

$$
\begin{aligned}
\nu_{-k}^{(-k)} &= (-1)^k \, \alpha_{-k+1} \, \alpha_k \, \beta_{k-1}^{-1} \, \delta_{k-1}^{-1} \, \nu_{-k+1}^{(-k+1)}, \\
\mu_{-k}^{(-k)} &= (-1)^k \, \alpha_{-k+1} \, \alpha_k \, \beta_{k-1}^{-1} \, \delta_k^{-1} \, \mu_{-k+1}^{(-k+1)}, \\
\nu_k^{(k)} &= (-1)^k \, \alpha_k^{-1} \, \alpha_{-k}^{-1} \, \beta_{k-1}^{-1} \, \delta_k^{-1} \, \nu_{k-1}^{(k-1)}, \\
\mu_{k+1}^{(k+1)} &= (-1)^k \, \alpha_k^{-1} \, \alpha_{-k}^{-1} \, \beta_k^{-1} \, \delta_k^{-1} \, \mu_k^{(k)}.
\end{aligned}
\tag{4.10}
$$

From these relations, the expressions for the coefficients follow easily. ∎

The results of Proposition 4.2 and Lemma 4.3 lead to the following corollary.

COROLLARY 4.4. *The polynomials p_k and p_{-k} can be expressed as*

$$
\begin{aligned}
p_k(t) &= \mu_k^{(k)} \cdot t^{-k+1} \cdot \det(tI_{2k-1} - R_{2k-1}), \\
p_{-k}(t) &= \mu_{-k}^{(-k)} \cdot t^k \cdot \det(t^{-1} I_{2k} - \widetilde{S}_{2k}),
\end{aligned}
$$

where $\mu_k^{(k)}$ and $\mu_{-k}^{(-k)}$ are defined in (4.8)*, and $\widetilde{S}_{2k}$ is the leading submatrix of order $2k$ of $\widetilde{S}_{2k+1}$ defined in* (3.5)*. The polynomials q_k and q_{-k} can be expressed analogously.*

We recall from Proposition 2.1 that for increasing k the largest singular value of H_{2k-1} converges monotonically to σ_1, and the inverse of the largest singular value of

K_{2k} converges monotonically to σ_n. This implies that the largest zero of polynomial p_k increases monotonically to σ_1^2. Likewise, the smallest zero of polynomial p_{-k} decreases monotonically to σ_n^2. These polynomials are used in the next section to obtain probabilistic bounds for both the largest and smallest singular value of A.

5. Probabilistic bounds for the condition number. After step k, extended Lanczos bidiagonalization implicitly provides Laurent polynomials p_k and p_{-k}. In the previous section we have seen that the zeros of p_k and p_{-k} are closely related to the singular values of the matrices H and K (Proposition 4.2). Moreover, the polynomials $|p_k|$ and $|p_{-k}|$ are strictly increasing to the right of their largest zero and also to the left of their smallest zero, for $t \to 0$. These properties will lead to the derivation of a probabilistic upper bound for $\kappa(A)$. Therefore, we first observe the two equalities

$$\begin{aligned} 1 = \|\mathbf{v}_k\|^2 &= \|p_k(A^TA)\mathbf{v}_0\|^2 &= \left\|\sum_{i=1}^n p_k(A^TA)\gamma_i\mathbf{y}_i\right\|^2 &= \sum_{i=1}^n \gamma_i^2 p_k(\sigma_i^2)^2, \\ 1 = \|\mathbf{v}_{-k}\|^2 &= \|p_{-k}(A^TA)\mathbf{v}_0\|^2 &= \left\|\sum_{i=1}^n p_{-k}(A^TA)\gamma_i\mathbf{y}_i\right\|^2 &= \sum_{i=1}^n \gamma_i^2 p_{-k}(\sigma_i^2)^2. \end{aligned}$$

Here we used, in view of (2.1), that $A^TA\mathbf{y}_i = \sigma_i^2\mathbf{y}_i$ and the fact that the right singular vectors $\mathbf{y}_i$ are orthonormal. Since the obtained sums only consist of nonnegative terms, we conclude that

$$|p_k(\sigma_1^2)| \le \frac{1}{|\gamma_1|} \qquad \text{and} \qquad |p_{-k}(\sigma_n^2)| \le \frac{1}{|\gamma_n|}. \tag{5.1}$$

Similarly,

$$\begin{aligned} 1 = \|\mathbf{u}_k\|^2 &= \|q_k(AA^T)A\mathbf{v}_0\|^2 &= \left\|\sum_{i=1}^n q_k(AA^T)\gamma_i\sigma_i\mathbf{x}_i\right\|^2 &= \sum_{i=1}^n \gamma_i^2\sigma_i^2 q_k(\sigma_i^2), \\ 1 = \|\mathbf{u}_{-k}\|^2 &= \|q_{-k}(AA^T)A\mathbf{v}_0\|^2 &= \left\|\sum_{i=1}^n q_{-k}(AA^T)\gamma_i\sigma_i\mathbf{x}_i\right\|^2 &= \sum_{i=1}^n \gamma_i^2\sigma_i^2 q_{-k}(\sigma_i^2). \end{aligned}$$

Here we used that $AA^T\mathbf{x}_i = \sigma_i^2\mathbf{x}_i$ and the fact that the left singular vectors $\mathbf{x}_i$ are orthonormal. Again, the sum we obtain only contains nonnegative terms and thus $1 \ge \sigma_1\,|\gamma_1|\,|q_k(\sigma_1^2)|$, which gives us the inequality

$$\sigma_1\,|q_k(\sigma_1^2)| \le \frac{1}{|\gamma_1|}, \qquad \sigma_n\,|q_{-k}(\sigma_n^2)| \le \frac{1}{|\gamma_n|}. \tag{5.2}$$

If γ_1 would be known, the first estimates in (5.1) and (5.2) would provide an upper bound for $\|A\|^2 = \sigma_1^2$, namely the largest zero of the functions

$$f_1(t) = |p_k(t)| - \frac{1}{|\gamma_1|}, \qquad f_2(t) = t\,|q_k(t)| - \frac{1}{|\gamma_1|}.$$

Similarly, if γ_n would be known, the second estimates in (5.1) and (5.2) would both provide a lower bound for $\|A^{-1}\|^{-2} = \sigma_n^2$, namely the smallest zero of the functions

$$g_1(t) = |p_{-k}(t)| - \frac{1}{|\gamma_n|}, \qquad g_2(t) = t\,|q_{-k}(t)| - \frac{1}{|\gamma_n|}.$$

However, both γ_1 and γ_n are unknown. Therefore, we will compute a value δ that will be a lower bound for $|\gamma_1|$ and $|\gamma_n|$ with a user-chosen probability. Suppose that $|\gamma_1| < \delta$. Then the largest zero of $f_1^\delta(t) = |p_k(t)| - \delta^{-1}$ is smaller than the largest zero of $f_1^{\gamma_1}(t) = |p_k(t)| - |\gamma_1|^{-1}$ and thus may be less then σ_1^2. This means that δ may not give an upper bound for σ_1. We now compute the value δ such that the probability that $|\gamma_1| < \delta$ (or $|\gamma_n| < \delta$) is small, namely ε. Let S^{n-1} be the unit sphere in $\mathbb{R}^n$. We choose the starting vector $\mathbf{v}_0$ randomly from a uniform distribution on S^{n-1} (MATLAB code: `v0=randn(n,1); v0=v0/norm(v0)`) (see, e.g., [17, p. 1116]), which (by an orthogonal transformation) implies that $(\gamma_1, \ldots, \gamma_n)$ is also random with respect to the uniform distribution on S^{n-1}.

LEMMA 5.1. *Assume that the starting vector $\mathbf{v}_0$ has been chosen randomly with respect to the uniform distribution over the unit sphere S^{n-1} and let $\delta \in [0,1]$. Then*

$$P(|\gamma_1| \le \delta) = 2B(\tfrac{n-1}{2}, \tfrac{1}{2})^{-1} \int_0^{\arcsin(\delta)} \cos^{n-2}(t)\, dt = B(\tfrac{n-1}{2}, \tfrac{1}{2})^{-1} \int_0^{\delta^2} t^{-\frac{1}{2}}(1-t)^{\frac{n-3}{2}}\, dt,$$

where B denotes Euler's Beta function: $B(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt$, and P stands for probability.

Proof. For the first equality, see [21, Lemma 3.1], and for the second, see [15, Theorem 7.1]. ◻

The user selects the probability $\varepsilon = P(|\gamma_1| \le \delta)$, i.e., the probability that the computed bound may not be an upper bound for the singular value σ_1. Given this user-chosen ε we have to determine the δ for which

$$\varepsilon = \frac{B_{\text{inc}}(\frac{n-1}{2}, \frac{1}{2}, \delta^2)}{B_{\text{inc}}(\frac{n-1}{2}, \frac{1}{2}, 1)}, \tag{5.3}$$

where the incomplete Beta function is defined as $B_{\text{inc}}(x,y,z) = \int_0^z t^{x-1}(1-t)^{y-1}dt$. The δ can be computed using MATLAB's function `betaincinv`. With this δ we can compute two probabilistic bounds, i.e., the square root of the largest zero of the function f_1^δ and the square root of the smallest zero of the function g_1^δ. Computing these values can be done with Newton's method or bisection. Note that one could equally choose to use the functions f_2^δ and g_2^δ. We thus acquire a probabilistic upper bound for σ_1 and a probabilistic lower bound for σ_n:

$$\sigma_1 < \sigma_{\text{up}}^{\text{prob}} \qquad \text{and} \qquad \sigma_n > \sigma_{\text{low}}^{\text{prob}}.$$

Both inequalities are true with a probability of at least $1-\varepsilon$. Since the coefficients γ_1 and γ_n are chosen independently, the probability that both inequalities hold is at least $1-2\varepsilon$. This proves the following theorem.

THEOREM 5.2. *Assume that the starting vector $\mathbf{v}_0$ has been chosen randomly with respect to the uniform distribution over S^{n-1}. Let $\varepsilon \in (0,1)$ and let δ be given by* (5.3). *Then $\sigma_{\text{up}}^{\text{prob}}$, the square root of the largest zero of the polynomial*

$$f_1^\delta(t) = |p_k(t)| - \tfrac{1}{\delta}, \tag{5.4}$$

is an upper bound for σ_1 with probability at least $1-\varepsilon$. Also, $\sigma_{\text{low}}^{\text{prob}}$, the square root of the smallest zero of the polynomial

$$g_1^\delta(t) = |p_{-k}(t)| - \tfrac{1}{\delta}, \tag{5.5}$$

is a lower bound for σ_n *with probability at least* $1-\varepsilon$.

Note that the implementation of the polynomial uses the recurrence relations in (4.9). Therefore, we approximate directly the singular values σ_1 and σ_n, avoiding taking squares or square roots. Combining these two bounds leads to a probabilistic upper bound for the condition number of A.

COROLLARY 5.3. *The inequality*

$$\kappa(A) = \frac{\sigma_1}{\sigma_n} \le \frac{\sigma_{\text{up}}^{\text{prob}}}{\sigma_{\text{low}}^{\text{prob}}} = \kappa_{\text{up}}(A) \tag{5.6}$$

holds with probability at least $1-2\varepsilon$.

The probabilistic upper bounds usually decrease monotonically as a function of k. The lemma below gives some intuition for this behavior.

LEMMA 5.4. *Let* t_1 *and* t_2 *be such that* $|p_k(t_1)| = \frac{1}{\delta}$, $|p_{k+1}(t_2)| = \frac{1}{\delta}$, *and define* $M := \alpha_k\alpha_{-k}\beta_k\delta_k$. *If* $t_1 \ge \theta_1^2 + M^{-1}(1+\sqrt{M}\theta_2)$, *then* $t_2 \le t_1$.

Proof. We investigate when $|p_{k+1}(t_1)| \ge \frac{1}{\delta}$, since this implies $t_2 \le t_1$. Denote by $\theta_1^2 \ge \cdots \ge \theta_{2k+1}^2$ the zeros of the polynomial $p_{k+1}(t)$, and by $\eta_1^2 \ge \cdots \ge \eta_{2k-1}^2$ the zeros of $p_k(t)$. Then

$$\left|\frac{p_{k+1}(t_1)}{p_k(t_1)}\right| = \left|\frac{\mu_{k+1}\, t_1^{-k}(t_1-\theta_1^2)\cdots(t_1-\theta_{2k+1}^2)}{\mu_k\, t_1^{-k+1}(t_1-\eta_1^2)\cdots(t_1-\eta_{2k-1}^2)}\right| = \delta\,|p_{k+1}(t_1)|.$$

The relations in (4.10) show that $|\frac{\mu_{k+1}}{\mu_k}| = (\alpha_k\alpha_{-k}\beta_k\delta_k)^{-1} =: M$. By the interlacing properties of singular values ($\eta_{2i-1} \ge \theta_{2i+1}$ for $i=1,\ldots,k$) we obtain the inequality

$$\delta\,|p_{k+1}(t_1)| \ge M\,\frac{(t_1-\theta_1^2)(t_1-\theta_2^2)}{t_1}.$$

So we are interested in finding t_1 such that $M\,(t_1-\theta_1^2)(t_1-\theta_2^2) \ge t_1$, which is

$$Mt_1^2 - (M(\theta_1^2+\theta_2^2)+1)t_1 + M\theta_1^2\theta_2^2 \ge 0.$$

This holds for

$$\begin{aligned} t_1 &\ge \tfrac{1}{2M}\Big(M(\theta_1^2+\theta_2^2)+1+\sqrt{(M(\theta_1^2+\theta_2^2)+1)^2-4M^2\theta_1^2\theta_2^2}\Big)\\ &= \tfrac{1}{2}(\theta_1^2+\theta_2^2)+\tfrac{1}{2M}+\tfrac{1}{2M}\sqrt{(M(\theta_1^2-\theta_2^2)+1)^2+4M\theta_2^2}. \end{aligned}$$

Therefore, $\delta\,|p_{k+1}(t_1)| \ge 1$ (and hence $t_2 < t_1$) holds for $t_1 \ge \theta_1^2 + M^{-1}(1+\sqrt{M}\theta_2)$. ∎

6. Other condition estimators. In this section we will first compare probabilistic results for $\kappa_2(A)$ obtained by Dixon [2] and Gudmundsson, Kenney, and Laub [6] with those of our method. Subsequently, we will briefly mention some condition number estimators for $\kappa_1(A)$ and $\kappa_F(A)$.

As for the method introduced in this paper, for all methods to approximate either $\kappa_1(A)$, or $\kappa_F(A)$, or $\kappa_2(A)$ discussed in this section, an LU decomposition is needed and $\mathcal{O}(1)$ vectors of storage are required (for our method see the recurrence relations (3.3) and the algorithm presented in section 7). Note that of the approaches discussed in this section only the block method by Higham and Tisseur [8] is also suitable for complex matrices.

6.1. Probabilistic condition estimators based on the 2-norm.

THEOREM 6.1 (see Dixon [2, Thm. 1]).[1] *Let B be a symmetric positive definite (SPD) matrix with eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$ and $\zeta > 1$. If $\mathbf{v}$ is chosen randomly on the unit sphere, then*

$$\mathbf{v}^T B \mathbf{v} \leq \lambda_1 \leq \zeta \cdot \mathbf{v}^T B \mathbf{v} \tag{6.1}$$

holds with probability at least $1 - 0.8\sqrt{n/\zeta}$.

Note that the left inequality always holds; the probabilistic part only concerns the second inequality. Dixon [2, Thm. 2] subsequently applies this result to both $B^k = (A^TA)^k$ and $B^{-k} = (A^TA)^{-k}$, which gives the following theorem.

THEOREM 6.2 (see Dixon [2, Thm. 2]). *Let A be a real nonsingular $n \times n$ matrix and let k be a positive integer. For $\mathbf{v}$, $\mathbf{w} \in \mathbb{R}^n$, define*

$$\varphi_k(\mathbf{v}, \mathbf{w}) = (\mathbf{v}^T (A^TA)^k \mathbf{v} \, \cdot \, \mathbf{w}^T (A^TA)^{-k} \mathbf{w})^{1/2k}.$$

If $\mathbf{v}$ and $\mathbf{w}$ are selected randomly and independently on S^{n-1} and $\zeta > 1$, then

$$\varphi_k(\mathbf{v}, \mathbf{w}) \leq \kappa(A) \leq \zeta \cdot \varphi_k(\mathbf{v}, \mathbf{w})$$

holds with probability at least $1 - 1.6\sqrt{n/\zeta^k}$.

Kuczyński and Woźniakowski [17] present several probabilistic bounds for quantities that are better estimates of the largest eigenvalue of an SPD matrix than the one considered by Dixon, with the same number of matrix-vector products. They appropriately call the method that leads to the quantity $(\mathbf{v}^T B^k \mathbf{v})^{1/k}$ studied by Dixon the *modified* power method. The more common power method considers, with the same number k of matrix-vector products, the Rayleigh quotient of $B^{k-1}\mathbf{v}$, that is, the quantity $(B^{k-1}\mathbf{v})^T B B^{k-1} \mathbf{v} = \mathbf{v}^T B^{2k-1} \mathbf{v}$. This generally results in a better approximation than the quantity considered by Dixon. In [17], the following results are given for the power method and the Lanczos method.

THEOREM 6.3 (see Kuczyński and Woźniakowski [17, Thm. 4.1(a)]). *With the same notation as in Theorem 6.1, let $0 < \eta < 1$. Let θ^{pow} be the largest Ritz value obtained with $k \geq 2$ steps of the power method. Then the probability that*

$$\lambda_1 < (1-\eta)^{-1}\, \theta^{\text{pow}} \tag{6.2}$$

holds is at least $1 - 0.824\sqrt{n}\,(1-\eta)^{k-\frac{1}{2}}$.

THEOREM 6.4 (see Kuczyński and Woźniakowski [17, Thm. 4.2(a)]). *With the same notation as in Theorem 6.3, let θ^{Lan} be the largest Ritz value obtained with k steps of Lanczos. Then the probability that*

$$\lambda_1 < (1-\eta)^{-1}\, \theta^{\text{Lan}} \tag{6.3}$$

holds is at least $1 - 1.648\sqrt{n}\, e^{-\sqrt{\eta}\,(2k-1)}$.

The proof of Theorem 6.4 uses a Chebyshev polynomial, a well-known proof technique in the area of Krylov methods. Extended Lanczos bidiagonalization adaptively constructs a polynomial that is optimal in some sense for the given matrix and starting vector. Therefore, as we will see below, our probabilistic bounds are usually better than that of Theorem 6.4.

[1] Note that [2, Thm. 1] contains a typo: k should be 1.

We can now apply Theorems 6.3 and 6.4 to $B = A^TA$ and $B = (A^TA)^{-1}$ as above. The following results are new, but follow directly from [17].

COROLLARY 6.5. *Let A be a real nonsingular $n \times n$ matrix and let k be a positive integer, and let $\mathbf{v}$ and $\mathbf{w}$ be random independent vectors on S^{n-1}.*

(a) [Power method on A^TA *and* $(A^TA)^{-1}$] *Let $\theta_{\max}^{\text{pow}} = \frac{\mathbf{v}^T(A^TA)^{2k-1}\mathbf{v}}{\mathbf{v}^T(A^TA)^{2k-2}\mathbf{v}}$ be the approximation to σ_1^2 obtained with k steps of the power method applied to A^TA with starting vector $\mathbf{v}$, and let $\theta_{\min}^{\text{pow}} = \frac{\mathbf{w}^T(A^TA)^{-(2k-1)}\mathbf{w}}{\mathbf{w}^T(A^TA)^{-(2k-2)}\mathbf{w}}$ be the approximation to σ_n^{-2} obtained with k steps of the power method applied to $(A^TA)^{-1}$ with starting vector $\mathbf{w}$. Then*

$$\kappa(A) \le (1-\eta)^{-1}\,(\theta_{\max}^{\text{pow}} \cdot \theta_{\min}^{\text{pow}})^{1/2}$$

holds with probability at least $1 - 1.648\,\sqrt{n}\,(1-\eta)^{k-\frac{1}{2}}$.

(b) [Lanczos on A^TA and $(A^TA)^{-1}$] *Let $\theta_{\max}^{\text{Lan}}$ be the largest Ritz value obtained with k steps of Lanczos applied to A^TA with starting vector $\mathbf{v}$, and let $\theta_{\min}^{\text{Lan}}$ be the largest Ritz value obtained with k steps of Lanczos applied to $(A^TA)^{-1}$ with starting vector $\mathbf{w}$. Then*

$$\kappa(A) \le (1-\eta)^{-1}\,(\theta_{\max}^{\text{Lan}} \cdot \theta_{\min}^{\text{Lan}})^{1/2}$$

holds with probability at least $1 - 3.296\,\sqrt{n}\,e^{-\sqrt{\eta}\,(2k-1)}$.

Example 6.6. We now give an indicative numerical example for the diagonal matrix $A =$ `diag(linspace(1,1e12,n))` of size $n = 10^5$ and $\kappa(A) = 10^{12}$. In Table 1, the probabilistic upper bounds by Dixon (the modified power method, Theorem 6.2), Kuczyński and Woźniakowski (the power method and the Lanczos method, Corollary 6.5), and the extended Lanczos bidiagonalization method are considered. We give the ratio $\kappa^{\text{up}}/\kappa^{\text{low}}(A)$, where κ^{up} denotes the various probabilistic upper bounds, where the requirement is that each holds with probability at least 98%. As expected, the power method gives a smaller ratio than the modified power method (see also [17] for more details). The ratio generated by a Chebyshev polynomial is even better, taking into account the subspace effect of a Krylov method. However, the ratio obtained with the polynomial implicitly generated by the method of this paper is the best.

TABLE 1

Ratios $\kappa^{\text{up}}(A)/\kappa^{\text{low}}(A)$ for $A =$ `diag(linspace(1,1e12,n))` of size $n = 10^5$, where κ^{up} denotes the probabilistic upper bound provided by Dixon [2], Kuczyński and Woźniakowski (Corollary 6.5), and the extended Lanczos bidiagonalization method (Ext LBD). We take $k = 10$, 20, and 30 steps, and require that all upper bounds κ^{up} hold with at least 98% ($\varepsilon = 0.01$).

k	Dixon	K&W (power)	K&W (Lanczos)	Ext LBD
10	7.60	2.92	1.49	1.16
20	2.76	1.68	1.08	1.04
30	1.97	1.41	1.04	1.02

6.2. Condition estimators based on other norms. Next, we mention the successful block method by Higham and Tisseur [8] to estimate the 1-norm condition number $\kappa_1(A) = \|A\|_1\,\|A^{-1}\|_1$. Although $\kappa_2(A)$ and $\kappa_1(A)$ are "equivalent" norms in $\mathbb{R}^n$ in the sense that $\frac{1}{n}\,\kappa_1(A) \le \kappa_2(A) \le n\,\kappa_1(A)$, these bounds are much too crude to be useful for large matrices. Therefore, we may well view $\kappa_2(A)$ and $\kappa_1(A)$ as independent quantities in practice; which one is preferred may depend on the user and application.

Gudmundsson, Kenney, and Laub [6] present an estimator for the condition number based on the Frobenius norm. They select k vectors from S^{n-1}, compute an orthonormal basis Q for the span, and take $\sqrt{n/k}\,\|AQ\|_F\,\|A^{-1}Q\|_F$ as an estimate for $\kappa_F(A)$. Again, although $\kappa_2(A)$ and $\kappa_F(A)$ are related in the sense that $\kappa_2(A) \le \kappa_F(A) \le n\,\kappa_2(A)$, they can be seen as independent quantities in practice.

7. Numerical experiments.

We present the pseudocode for the extended Lanczos bidiagonalization method including the computation of a lower bound and a probabilistic upper bound for the condition number. This pseudocode shows that this method requires only three vectors of storage. Because of the modest number of steps needed to achieve the given ratio, it turns out that in our examples reorthogonalization with respect to more previous vectors is not needed.

Algorithm: Extended Lanczos bidiagonalization method with lower and probabilistic upper bounds.

Input: Nonsingular $(n \times n)$ matrix A, random starting vector $\mathbf{w} = \mathbf{v}_0$, probability level ε, ratio ζ, maximum extended Krylov dimension $2k$.

Output: A lower bound $\kappa_{\text{low}}(A)$ and a probabilistic upper bound $\kappa_{\text{up}}(A)$ for the condition number $\kappa(A)$ such that $\kappa_{\text{up}}/\kappa_{\text{low}} \le \zeta$. The probability that $\kappa(A) \le \kappa_{\text{up}}(A)$ holds is at least $1 - 2\varepsilon$. In the unlikely event of a breakdown, the algorithm aborts and may not return any estimate.

```
Determine δ from n and ε, see (5.3).
for j = 0, ..., k−1
    u = A w
    α_{−j} = ‖u‖
    if α_{−j} = 0, abort, end
    u = u / α_{−j}
    u = A^T u
    if j > 0
        β_{−j} = v^T u
        u = u − β_{−j} v
    end
    u = u − α_{−j} w
    β_j = ‖u‖
    if β_j = 0, abort, end
    v = u / β_j
    u = A^{−T} v
    if ‖u‖ = 0, abort, end
    α_{j+1} = ‖u‖^{−1}
    Create H_{2(j+1)} using the obtained coefficients α's and β's (see (3.1)).
    Determine largest and smallest singular values θ_1 and θ_{2(j+1)}, respectively, of H_{2(j+1)}.
    Compute lower bound κ_low(A) = θ_1/θ_{2(j+1)} for κ(A) (see (3.4)).
    Determine σ_up^prob for σ_1 with probability ≥ 1 − ε using f_1^δ (see (5.4)).
    Determine σ_low^prob for σ_n with probability ≥ 1 − ε using g_1^δ (see (5.5)).
    Compute probabilistic upper bound κ_up(A) = σ_up^prob / σ_low^prob for κ(A) (see (5.6)).
    if κ_up/κ_low ≤ ζ, quit, end
    u = α_{j+1} u
    u = A^{−1} u
    δ_{−j} = w^T u
    u = u − δ_{−j} w − α_{j+1}^{−1} v
    δ_{j+1} = ‖u‖
    if δ_{j+1} = 0, abort, end
    w = u / δ_{j+1}
end
```

Experiment 7.1. First, we test the method on some well-known diagonal test matrices to get an impression of the performance of the method. In Figure 1, we plot the convergence of the probabilistic upper bound $\kappa_{\rm up}(A)$ and lower bound $\kappa_{\rm low}(A)$ as a function of k for the matrix $A =$ `diag(linspace(1,1e12,n))`, for $n = 10^5$ (a) and for an "exponential diagonal" matrix of the form $A =$ `diag(`ρ`.^(0:1e5-1))`, where ρ is such that $\kappa(A) = 10^{12}$ (b). The plots suggest that the spectrum of the latter matrix is harder.

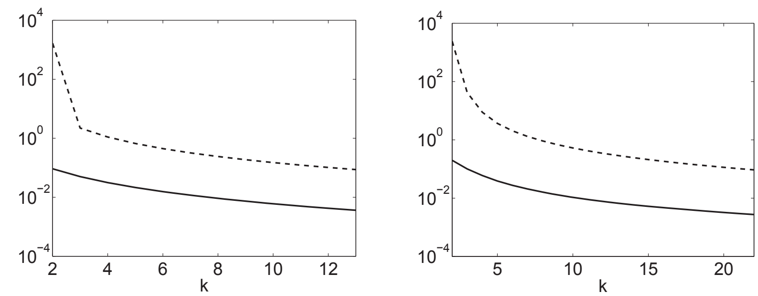

(a) (b)

Fig. 1. *The relative errors* $\kappa_{\rm up}(A)/\kappa(A) - 1$ *(dash) and* $1 - \kappa_{\rm low}(A)/\kappa(A)$ *(solid) as a function of* k*, for* $A =$ *diag(linspace(1,1e12,n)),* $n = 10^5$ (a)*, and a matrix of the form* $A =$ diag(ρ.^(0:1e5-1)) *with* $\kappa(A) = 10^{12}$ (b). *Here,* $\kappa_{\rm low}(A)$ *is a lower bound and* $\kappa_{\rm up}(A)$ *is an upper bound with probability at least* 98%.

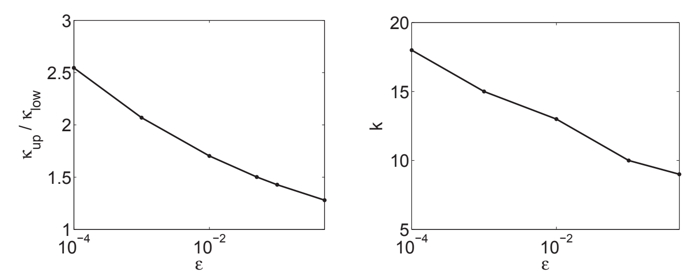

(a) (b)

Fig. 2. *For* $A =$ *diag(linspace(1,1e12,n)),* $n = 10^5$*, after* $k = 5$ *steps of the method:* (a) *the ratio* $\kappa_{\rm up}(A)/\kappa_{\rm low}(A)$*, where* $\kappa_{\rm up}(A)$ *is an upper bound with probability at least* $1 - 2\varepsilon$*, as a function of* ε*;* (b) *the iteration* k *needed to ensure that* $\kappa_{\rm up}(A) \le 1.1 \cdot \kappa_{\rm low}(A)$*, where* $\kappa_{\rm up}(A)$ *is an upper bound with probability at least* $1 - 2\varepsilon$*, as a function of* ε.

Next, for Figure 2(a), we carry out $k = 5$ steps of the method for $A =$ `diag(linspace(1,1e12,n))`, $n = 10^5$, and investigate the behavior of the ratio $\kappa_{\rm up}(A)/\kappa_{\rm low}(A)$, where $\kappa_{\rm up}(A)$ is an upper bound with probability at least $1 - 2\varepsilon$, as a function of ε. In Figure 2(b) we plot the iteration k that is needed to ensure that $\kappa_{\rm up}(A) \le 1.1 \cdot \kappa_{\rm low}(A)$, as a function of ε.

Experiment 7.2. Next, we test the method to estimate the condition number for some large matrices. The matrices we choose are real and nonsymmetric. Most of these matrices can be found in the Matrix Market [18] or the University of Florida Sparse Matrix Collection [1, 4]. The starting vector $\mathbf{v}_0$ is randomly chosen from a uniform distribution on S^{n-1} as explained in section 5. For these experiments we choose $\varepsilon = 0.01$, which corresponds to a reliability of at least 98% for the bounds for the condition number to be true (see section 5). Also we choose $\zeta = 2$ and $\zeta = 1.1$ such that the ratio of the probabilistic upper bound and the lower bound is $\leq \zeta$. To accomplish this, the method adaptively chooses the a number of steps k. Note that k steps correspond to k operations with A^TA and k operations with $(A^TA)^{-1}$. We use MATLAB's betaincinv to compute δ and bisection to compute the largest and smallest zero of f_1^δ and g_1^δ, respectively, (see (5.4) and (5.5)).

In Table 2 the results for $\zeta = 2$ are presented. The reason for the choice of $\zeta = 2$ is a comparison of our method to the block method by Higham and Tisseur [8] to estimate the 1-norm condition number $\kappa_1(A)$, which is reported to give almost always an estimate correct to within a factor of 2. Although κ_1 and κ_2 are independent quantities (see section 6.2 for comments), the methods have both a storage of $\mathcal{O}(1)$ vectors and for both methods (only) one LU-factorization is computed which is needed for the inverse operations A^{-1} and A^{-T}. The comparison is made to indicate that the running time of the two methods usually does not differ much (see Table 2). As is shown in Table 2, especially for the larger matrices, a large part of the computational time is spent on the computation of the LU-factorization. Therefore, for such matrices extended Lanczos bidiagonalization may be seen as a relatively cheap add-on. For $\zeta = 2$, usually only a modest number of steps k are sufficient. Of course, choosing a larger ζ will decrease this number of steps even more. While decreasing ζ will make the method computationally more expensive, for many matrices this will be a relatively small increase in view of the costs of the LU decomposition. In Table 3 the results for $\zeta = 1.1$, giving very sharp bounds, show that even for this small ζ the number of steps k and the running time remain modest.

Table 2

The approximations of the condition number κ of different matrices using extended Lanczos bidiagonalization. The method gives a lower bound κ_{low} for κ and also a probabilistic upper bound κ_{up} that holds with probability at least 98% ($\varepsilon = 0.01$). The method continues until the ratio $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$ is below the indicated level of $\zeta = 2$. The number of steps k needed to obtain this ratio and the CPU time in seconds are shown. Also the percentage of the time taken by the LU decomposition is displayed. Lastly, we give CPU[1] of `condest(A)`*. The symbol $\star$ is used when the value is too expensive to compute.*

Matrix A	Dim.	κ	κ_{low}	κ_{up}	k	CPU	LU	CPU[1]
utm5940	5940	$4.35 \cdot 10^8$	$3.98 \cdot 10^8$	$7.21 \cdot 10^8$	4	0.17	55	0.12
dw8192	8192	$3.81 \cdot 10^6$	$3.81 \cdot 10^6$	$5.07 \cdot 10^6$	6	0.10	49	0.08
grcar10000	10000	$3.63 \cdot 10^0$	$3.59 \cdot 10^0$	$5.80 \cdot 10^0$	6	0.07	30	0.05
memplus	17758	$1.29 \cdot 10^5$	$1.29 \cdot 10^5$	$2.47 \cdot 10^5$	6	0.17	57	0.13
af23560	23560	$1.99 \cdot 10^4$	$1.93 \cdot 10^4$	$2.82 \cdot 10^4$	6	0.93	73	0.88
rajat16	96294	$\star$	$5.63 \cdot 10^{12}$	$5.69 \cdot 10^{12}$	5	9.29	97	9.19
torso1	116158	$\star$	$1.41 \cdot 10^{10}$	$1.42 \cdot 10^{10}$	3	26.8	94	28.5
dc1	116835	$\star$	$2.39 \cdot 10^8$	$4.59 \cdot 10^8$	5	5.87	94	5.57
twotone	120750	$\star$	$1.75 \cdot 10^9$	$2.91 \cdot 10^9$	4	1.37	75	1.33
FEM-3D-thermal2	147900	$\star$	$3.05 \cdot 10^3$	$5.15 \cdot 10^3$	7	13.1	80	12.7
xenon2	157464	$\star$	$4.29 \cdot 10^4$	$8.14 \cdot 10^4$	7	19.4	83	19.6
crashbasis	160000	$\star$	$6.30 \cdot 10^2$	$1.21 \cdot 10^3$	7	3.35	64	2.59
scircuit	170998	$\star$	$2.40 \cdot 10^9$	$4.69 \cdot 10^9$	7	2.11	58	1.39
transient	178866	$\star$	$1.02 \cdot 10^{11}$	$2.00 \cdot 10^{11}$	8	7.76	87	7.12
stomach	213360	$\star$	$4.62 \cdot 10^1$	$9.02 \cdot 10^1$	6	13.7	80	13.7

Experiment 7.3. We compare the new method with the following alternative method to derive a lower bound for $\kappa_2(A)$. First, one applies k Lanczos iterations with A^TA to a starting vector $\mathbf{v}$, providing an approximation to $\sigma_1(A)$ from the standard Krylov subspace $\mathcal{K}_{k+1}(A^TA, \mathbf{v}) = \mathcal{K}^{1,k+1}(A^TA, \mathbf{v})$. Subsequently, one applies k Lanczos iterations with $(A^TA)^{-1}$ to the same starting vector $\mathbf{v}$, giving an approximation

Table 3

The approximations of the condition number κ of different matrices using extended Lanczos bidiagonalization. The method gives a lower bound κ_{low} for κ and also a probabilistic upper bound κ_{up} that holds with probability at least 98% ($\varepsilon = 0.01$). The method continues until the ratio $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$ is below the indicated level of $\zeta = 1.1$. The number of steps k needed to obtain this ratio and the CPU time in seconds are shown. Also, the percentage of the time taken by the LU decomposition is displayed. The symbol $\star$ is used when the value is too expensive to compute.

Matrix A	Dim.	κ	κ_{low}	κ_{up}	k	CPU	LU
utm5940	5940	$4.35 \cdot 10^8$	$4.35 \cdot 10^8$	$4.71 \cdot 10^8$	10	0.36	60
dw8192	8192	$3.81 \cdot 10^6$	$3.81 \cdot 10^6$	$4.06 \cdot 10^6$	9	0.13	35
grcar10000	10000	$3.63 \cdot 10^0$	$3.62 \cdot 10^0$	$3.97 \cdot 10^0$	13	0.13	20
memplus	17758	$1.29 \cdot 10^5$	$1.29 \cdot 10^5$	$1.41 \cdot 10^5$	15	0.28	31
af23560	23560	$1.99 \cdot 10^4$	$1.99 \cdot 10^4$	$2.12 \cdot 10^4$	9	1.10	63
rajat16	96294	$\star$	$5.63 \cdot 10^{12}$	$5.69 \cdot 10^{12}$	5	9.39	97
torso1	116158	$\star$	$1.41 \cdot 10^{10}$	$1.42 \cdot 10^{10}$	3	26.9	94
dc1	116835	$\star$	$2.39 \cdot 10^8$	$2.45 \cdot 10^8$	8	6.01	92
twotone	120750	$\star$	$1.75 \cdot 10^9$	$1.91 \cdot 10^9$	7	1.69	64
FEM-3D-thermal2	147900	$\star$	$3.15 \cdot 10^3$	$3.43 \cdot 10^3$	12	15.1	70
xenon2	157464	$\star$	$4.32 \cdot 10^4$	$4.67 \cdot 10^4$	14	22.5	71
crashbasis	160000	$\star$	$6.40 \cdot 10^2$	$7.01 \cdot 10^2$	18	5.21	40
scircuit	170998	$\star$	$2.45 \cdot 10^9$	$2.67 \cdot 10^9$	16	3.29	37
transient	178866	$\star$	$1.03 \cdot 10^{11}$	$1.11 \cdot 10^{11}$	21	9.24	73
stomach	213360	$\star$	$4.82 \cdot 10^1$	$5.24 \cdot 10^1$	14	17.3	64

to $\sigma_n(A)$ from the subspace $\mathcal{K}_{k+1}((A^TA)^{-1}, \mathbf{v}) = \mathcal{K}^{k+1,1}(A^TA, \mathbf{v})$. Together these two values form a lower bound for $\kappa(A)$ as in (3.4). The lower bound of extended Lanczos bidiagonalization is always at least as good as the lower bound obtained by the alternative approach, as the former approach considers subspaces of the extended space $\mathcal{K}^{k+1,k+1}(A^TA, \mathbf{v})$. Furthermore, since in the extended Lanczos bidiagonalization procedure we can control the ratio ζ, a natural stopping criterion arises for this method, as well as a good measure of the quality of both upper and lower bound. For the other approach these features are both missing.

As an example, the lower bound of $\kappa(A)$ for the matrix $A =$ `af23560` using extended Lanczos bidiagonalization ($k = 6$) is $1.93 \cdot 10^4$ in 0.93 seconds. Using twice a Lanczos procedure ($k = 6$) gives the lower bound $1.87 \cdot 10^4$ in 0.99 seconds. For the same number of steps, the matrix `stomach` gives $4.62 \cdot 10^1$ for extended Lanczos bidiagonalization (13.7 seconds) and $4.54 \cdot 10^1$ for the alternative approach (14.5 seconds). Besides a better lower bound, an important advantage of extended Lanczos bidiagonalization is that, almost for free, a probabilistic upper bound is provided as well. Note that in this example the CPU time for extended Lanczos bidiagonalization includes the time for the computation of the probabilistic upper bounds.

Experiment 7.4. Another alternative to approximate the condition number of A is to use the `svds` command in MATLAB. We compared our method, with the parameters $\zeta = 1.1$ and $\varepsilon = 0.01$, to the outcome of the command `svds(A,1,'L')/svds(A,1,0)`. The results in Table 4 show that our method significantly outperforms the `svds` approach concerning the running time (in these examples our method is 8 to 13 times faster), giving the same lower bound for $\kappa(A)$. Again, as stated in the previous experiment, our method also gives a probabilistic upper bound for the condition number almost for free.

Finally, some words on a breakdown. A breakdown takes place when the method has found an invariant subspace. This is a rare event; in exact arithmetic the probability that this happens for a $k \ll n$ is zero since we have selected a random vector.

Table 4

For three matrices the lower bound given by the extended Lanczos bidiagonalization (Ext LBD), with the parameters $\zeta = 1.1$ and $\varepsilon = 0.01$, and the corresponding CPU time in seconds is shown. Also, the bound given by the procedure using `svds` *and the corresponding CPU time in seconds is shown.*

Matrix	Ext LBD	CPU	svds	CPU
`memplus`	$1.294 \cdot 10^5$	0.28	$1.294 \cdot 10^5$	2.30
`af23560`	$1.988 \cdot 10^4$	1.10	$1.989 \cdot 10^4$	15.1
`rajat16`	$5.629 \cdot 10^{12}$	9.39	$5.629 \cdot 10^{12}$	100.9

A breakdown has not been encountered in our numerical experiments. However, it might happen in rare cases that in the algorithm α_{-j} (Step 5), β_j (Step 13), $\|\mathbf{u}\|$ (Step 15), or δ_{j+1} (Step 28) are zero or very small. In such a case, we can just stop the method and return the lower and probabilistic bounds obtained before the breakdown. If these do not yet satisfy the requirements of the user, we can restart the method with a new random vector. An extra run of the extended Lanczos bidiagonalization method will not increase the overall costs by much. With this adaptation we trust that the method can result in a robust implementation for the use in libraries.

8. Discussion and conclusions. We have proposed a new extended Lanczos bidiagonalization method. This method leads to tridiagonal matrices with a special structure. The method provides a lower bound for $\kappa(A)$ of good quality and a probabilistic upper bound for $\kappa(A)$ that holds with a user-chosen probability $1 - 2\varepsilon$. Although we have not encountered any breakdown in the experiments, the algorithm may abort and not return any estimate. When choosing k adaptively, given a user-selected ε and desired ratio $\kappa_{\mathrm{up}}(A)/\kappa_{\mathrm{low}}(A) < \zeta$, the results show that generally this k is fairly small, even for $\zeta = 1.1$. Only three vectors of storage are required. This method can be used whenever an LU-factorization is computable in a reasonable amount of time. (When this is not an option, methods such as the one in [9, 10] can be used.)

Acknowledgments. We would like to thank Lothar Reichel for fruitful conversations, and the two referees for very useful comments, which led to considerable improvements of this paper.

REFERENCES

[1] T. A. Davis and Y. Hu, *The University of Florida sparse matrix collection*, ACM Trans. Math. Software, 38 (2011), article 1.

[2] J. D. Dixon, *Estimating extremal eigenvalues and condition numbers of matrices*, SIAM J. Numer. Anal., 20 (1983), pp. 812–814.

[3] V. L. Druskin and L. A. Knizhnerman, *Extended Krylov subspaces: Approximation of the matrix square root and related functions*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 755–771.

[4] The University of Florida Sparse Matrix Collection, *A Repository for Test Matrices*, www.cise.ufl.edu/research/sparse/matrices/.

[5] G. Golub and W. Kahan, *Calculating the singular values and pseudo-inverse of a matrix*, J. Soc. Indust. Appl. Math. Ser. B Numer. Anal., 2 (1965), pp. 205–224.

[6] T. Gudmundsson, C. S. Kenney, and A. J. Laub, *Small-sample statistical estimates for matrix norms*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 776–792.

[7] N. Halko, P. G. Martinsson, and J. A. Tropp, *Finding structures with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.

[8] N. J. Higham and F. Tisseur, *A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1185–1201.
[9] M. E. Hochstenbach, *A Jacobi–Davidson type SVD method*, SIAM J. Sci. Comput., 23 (2001), pp. 606–628.
[10] M. E. Hochstenbach, *Harmonic and refined extraction methods for the singular value problem, with applications in least squares problems*, BIT, 44 (2004), pp. 721–754.
[11] M. E. Hochstenbach, *Probabilistic upper bounds for the matrix two-norm*, J. Sci. Comput., 57 (2013), pp. 464–476.
[12] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1991.
[13] C. Jagels and L. Reichel, *The extended Krylov subspace method and orthogonal Laurent polynomials*, Linear Algebra Appl., 431 (2009), pp. 441–458.
[14] C. Jagels and L. Reichel, *Recursion relations for the extended Krylov subspace method*, Linear Algebra Appl., 434 (2011), pp. 1716–1732.
[15] E. R. Jessup and I. C. F. Ipsen, *Improving the accuracy of inverse iteration*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 550–572.
[16] L. Knizhnerman and V. Simoncini, *A new investigation of the extended Krylov subspace method for matrix function evaluations*, Numer. Linear Algebra Appl., 17 (2010), pp. 615–638.
[17] J. Kuczyński and H. Woźniakowski, *Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1094–1122.
[18] The Matrix Market, math.nist.gov/MatrixMarket, a repository for test matrices.
[19] B. N. Parlett, *The Symmetric Eigenvalue Problem*, Classics Appl. Math. 20 SIAM, Philadelphia, 1998.
[20] V. Simoncini, *A new iterative method for solving large-scale Lyapunov matrix equations*, SIAM J. Sci. Comput., 29 (2007), pp. 1268–1288.
[21] J. L. M. van Dorsselaer, M. E. Hochstenbach, and H. A. van der Vorst, *Computing probabilistic bounds for extreme eigenvalues of symmetric matrices with the Lanczos method*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 837–852.